arXiv:0902.0342v1 [math.ST] 2 Feb 2009

ON A CONJECTURED SHARPNESS PRINCIPLE FOR PROBABILISTIC FORECASTING WITH CALIBRATION

SOUMIK PAL

Abstract. This note proves a weak type of the sharpness principle as conjectured by Gneiting, Balabdaoui, and Raftery [9] in connection with probabilistic forecasting subject to calibration constraints. A strong version of such a principle still awaits a proper formulation.

1. Introduction

This note is concerned with the conjectured sharpness principle described in the Gneiting-Balabdaoui-Raftery article on probabilistic forecasts [9] subject to calibration constraints. Using predictive distributions for making forecasts has been steadily gaining acceptance in various fields such as meteorology (Gneiting & Raftery [10]), economics (Granger [12]), and finance (Duffie & Pan [3]). The conceptual appeal of probabilistic forecasts and the advances in Markov Chain Monte Carlo methodology have made them very attractive. For a recent editorial on probabilistic forecasting, please see the article by Gneiting [8].

Predictive performances of forecasts have been traditionally measured by what is known as the probability integral transform or PIT. The PIT essentially involves applying the forecasted distribution functions to their corresponding true observations and checking the uniformity of the resulting histogram. This method first proposed by Dawid [4] and further by Diebold, Gunther, & Tay [7], only leads to a necessary condition for a forecaster to be ideal. In fact, that using the PIT just by itself can lead to erroneous conclusions was succinctly demonstrated by an example of Hamill [13]. To address this issue Gneiting et al. in [9] introduce the concept of maximizing the sharpness of the predictive distributions subject to calibrations. They define calibration as 'the statistical consistency between the predictive distributions and the associated observations, and is a joint property of the predictions and the values that materialize'. Sharpness, on the other hand, 'refers to the concentration of the predictive distributions and is a property of the forecasts only'. The authors then define a theoretical framework to assess calibration and sharpness and distinguish between several modes of calibration. Our article is based on this framework.

Let $\{F_t,\ t = 1, 2, \ldots\}$ and $\{G_t,\ t = 1, 2, \ldots\}$ denote sequences of continuous and strictly increasing distribution functions on the real line, possibly depending on stochastic parameters. One thinks of the G_t's as the true data-generating distributions and the corresponding sequence of F_t's as the sequence of probabilistic forecasts. Among all the modes of calibration proposed in [9], probably the most useful notion is that of probabilistic calibration. The sequence $\{F_t\}$ is probabilistically calibrated

Date: February 2, 2009.

relative to the sequence $\{G_t\}$ if

$$\lim_{T\to\infty} \frac{1}{T}\sum_{i=1}^{T} G_i \circ F_i^{-1}(p) = p, \quad \forall\, p \in (0,1). \tag{1}$$

It is not difficult to see that probabilistic calibration is essentially equivalent to the uniformity of the PIT histogram, a fact that leads to the practicality of its applications.

We should mention that calibration in the context of forecasting binary events is well studied in Statistics literature. See, for example, the articles by DeGroot & Fienberg [6], Dawid & Vovk [5], and Shafer & Vovk [14]. The elegant game-theoretic approach of Vovk & Shafer [15] is a natural culmination of this approach. The models considered here, however, are continuous and not within the purview of the previous approach.

2. Main results

The particular problem addressed in this note is the one stated in Section 2.4 of [9]. It involves what the authors of [9] call the sharpness principle. Ideally one would like to forecast the data generating distribution perfectly. Supposing that one maximizes sharpness among predictive distributions subject to calibration, how close to perfect forecasts can she get ? The strong sharpness principle in [9] states that perfect forecasts and maximization of sharpness subject to calibration is indeed the same. For probabilistic calibration it is difficult to find natural examples where such a principle holds. Thus, in [9], the authors also formulate the weak sharpness principle. Suppose the sequence $\{F_t\}$ is probabilistically calibrated with respect to $\{G_t\}$. If one assesses sharpness by variance, what can one deduce about the accuracy of the forecasts F_t's ? A partial answer, in the case of climatological forecaster (when all the F_t's are the same), is provided in Theorem 1 of [9]. It is shown that on average the forecasts will have a bigger variance than the actual observations. A similar conclusion for the general case is conjectured, and this is what we settle here. The choice of averaged variance as a measure of sharpness is somewhat arbitrary. Please see section 3 in Gneiting et al. [11] for further discussion.

The main result is the following proposition.

Proposition 1. *Suppose, for some $T \in \mathbb{N}$, the sequence of probability distributions $\{F_i\}$ and $\{G_i\}$ satisfy the finite probabilistic calibration condition*

$$\frac{1}{T}\sum_{1}^{T} G_t \circ F_t^{-1}(p) = p, \qquad \text{for all } p \in (0,1). \tag{2}$$

For all such sequences $(F_1, F_2, \ldots, F_T)$ and $(G_1, G_2, \ldots, G_T)$ satisfying (2), we have

$$\frac{1}{T}\sum_{1}^{T} Var(F_i) \geq \frac{1}{T}\sum_{1}^{T} Var(G_i), \tag{3}$$

with equality if and only if

$$E(G_1) - E(F_1) = E(G_2) - E(F_2) = \ldots = E(G_T) - E(F_T). \tag{4}$$

Here $E(H)$ and $Var(H)$ denote the expectation and the variance of a random variable with distribution function H.

With stronger assumptions one can extend the finite horizon result to a limiting result consistent with the definition of probabilistic calibration given in [9]. Below we outline one such approach. Our assumptions are deliberately kept stronger than required in order to maintain simplicity of the statement.

Proposition 2. *Assume that there is a bounded interval (a, b) such that each F_i is zero to the left of a and one to the right of b. Further assume that the sequences $\{F_i\}$ and $\{G_i\}$ satisfy the probabilistic calibration condition* (1) *and that there is a function $\theta : (0, 1) \to \mathbb{R}$ such that*

$$\lim_{T\to\infty} \sup_{u\in(0,1)} \left| \frac{1}{T} \sum_{i=1}^{T} \left(F_i^{-1}(u)\right)^2 - \theta(u) \right| = 0. \tag{5}$$

Then

$$\liminf_{T\to\infty} \frac{1}{T} \sum_{i=1}^{T} Var(F_i) \geq \limsup_{T\to\infty} \frac{1}{T} \sum_{i=1}^{T} Var(G_i). \tag{6}$$

Here lim sup *and* lim inf *denote the limit superior and the limit inferior of a real valued sequence.*

The significance of the previous inequalities is directly linked to the Gneiting-Balabdaoui-Raftery [9] paradigm of maximizing sharpness of a forecast subject to calibration. In practice, competing probabilistic forecasters need to be assessed on the basis of their predictive distributions and the outcomes that materialize, in accordance with Dawid's [4] prequential principle. Ideally we would like a forecaster to predict the true distribution, i.e. $F_t = G_t$ for all times t. However, G_t is unobservable, so this is not an operational criterion. Instead [9] proposes maximizing sharpness of the predictive distributions subject to calibration. Our results lend further support to their proposal, in the sense that if we measure sharpness by average variance, a probabilistically calibrated forecaster will not be sharper than the ideal forecaster.

However, Proposition 1 allows equality in (3) under apparently weaker conditions (4) than distributional equality. The quest for notions of calibration that lead to the ideal forecaster as the unique sharpest forecaster (a strong form of the sharpness principle) remains open.

Let us add that one should compare these results to existing results about decomposition of strictly proper scores into sharpness and reliability like terms for both binary forecasting and more general set-up. Please see the recent article by Bröcker [2] for a complete result and references of previous work. Such a decomposition, of course, will be the most complete version of a sharpness principle that one can desire. However, in our framework, it is not obvious how to obtain a similar decomposition theorem.

The rest of this note contains the proofs of the stated results.

3. Proofs

Proof of Theorem 1. Let $Z = (Z_1, Z_2, \ldots, Z_T)$ be a multinomial random vector with an equal probability of success in each of the T bins. Let U be the random variable

$$U := F_1(X_1)^{Z_1} F_2(X_2)^{Z_2} \ldots F_T(X_T)^{Z_T}$$

where $X = (X_1, X_2, \ldots, X_T)$ is a random vector such that marginally $X_1 \sim G_1, \ldots, X_T \sim G_T$ and is independent of Z. Then, as was mentioned in the paper [9], the finite probabilistic calibration condition implies U is a Uni$(0,1)$ random variable.

For (u, z) among the possible values of (U, Z), define the function

$$H(z,u) = \left(F_1^{z_1} F_2^{z_2} \ldots F_T^{z_T}\right)^{-1}(u). \tag{7}$$

We are going to compute $Var(H(Z,U))$ in two different methods.

The first method is to condition on $Z = z$, and use the decomposition

$$Var(H(Z,U)) = E\, Var(H|Z=z) + Var\, E(H|Z=z). \tag{8}$$

Now, it is straightforward to see that

$$\begin{aligned} Var\left(H|Z=z\right) &= Var(X_i), \quad \text{if } z_i = 1, z_j = 0\ j \neq i \\ E\left(H|Z=z\right) &= E(X_i), \quad \text{if } z_i = 1, z_j = 0\ j \neq i. \end{aligned} \tag{9}$$

Thus,

$$Var(H(Z,U)) = \frac{1}{T}\sum_{i=1}^{T} Var(X_i) + \frac{1}{T}\sum_{i=1}^{T}(\mu_i - \bar{\mu})^2.$$

Here $\mu_i = E(G_i)$ and $\bar{\mu}$ is the mean of the μ_i's. Thus

$$Var(H(Z,U)) \geq \frac{1}{T}\sum_{i=1}^{T} Var(G_i) \tag{10}$$

with equality if and only if all the G_i's have the same mean.

The second method to compute the variance of $H(Z,U)$ is to first condition on U, and repeat the variance decomposition formula

$$Var(H(Z,U)) = E\, Var(H|U=p) + Var\, E(H|U=p). \tag{11}$$

To compute the above, note that

$$\begin{aligned} E(H|U=p) = \frac{1}{T}\sum_{i=1}^{T} F_i^{-1}(p), \quad E(H^2|U=p) &= \frac{1}{T}\sum_{i=1}^{T}\left(F_i^{-1}(p)\right)^2, \\ Var(H|U=p) = \frac{1}{T}\sum_{i=1}^{T}\left(F_i^{-1}(p)\right)^2 - \frac{1}{T^2}&\left(\sum_{i=1}^{T} F_i^{-1}(p)\right)^2. \end{aligned} \tag{12}$$

Define $Y_i = F_i^{-1}(U)$, $i = 1, 2, \ldots, T$. By the finite probabilistic calibration condition, U is marginally uniform. Hence, each Y_i has distribution F_i. Now, from the equalities in (11) and (12), we get

$$\begin{aligned} Var(H(Z,U)) &= \frac{1}{T}\sum_{i=1}^{T} E\left(Y_i^2\right) - \frac{1}{T^2} E\left(\sum_{i=1}^{T} Y_i\right)^2 \\ &\quad + \frac{1}{T^2} Var\left(\sum_{i=1}^{T} Y_i\right) \\ &= \frac{1}{T}\sum_{i=1}^{T} E\left(Y_i^2\right) - \frac{1}{T^2}\left(\sum_{i=1}^{T} E(Y_i)\right)^2. \end{aligned} \tag{13}$$

The finally equality follows from the fact that for any Z, one has $Var(Z) = EZ^2 - (EZ)^2$.

Thus, we get the equality

$$\begin{aligned} Var(H(Z,U)) &= \frac{1}{T}\sum_{i=1}^{T} Var(Y_i) + \frac{1}{T}\sum_{1}^{T}(EY_i)^2 - \frac{1}{T^2}\left(\sum_{i=1}^{T} E(Y_i)\right)^2, \\ &= \frac{1}{T}\sum_{i=1}^{T} Var(Y_i) + \frac{1}{T}\sum_{i=1}^{T}(\alpha_i - \bar{\alpha})^2, \end{aligned} \tag{14}$$

where $\alpha_i = E(F_i)$ and $\bar{\alpha}$ is the mean of the α_i's.

Thus, if all α_i's (i.e., all the means of F_i's) are the same, we get exact equality above. In that case, combining with inequality (10), we get the result

$$\frac{1}{T}\sum_{i=1}^{T} Var(F_i) \geq \frac{1}{T}\sum_{i=1}^{T} Var(G_i), \tag{15}$$

with equality if and only if the means of all the G_i's are the same.

To prove (3), in case when all the α_i's are not the same, define

$$F_i^*(x) = F_i(x + \alpha_i), \quad G_i^*(x) = G_i(x + \alpha_i). \tag{16}$$

Notice, that the calibration condition (2) continues to hold. Additionally, each F_i^* has now mean zero and the same variance as F_i. Applying the variance inequality (15) for case of $\{F_i^*\}$ and $\{G_i^*\}$, we get the complete result as in the statement. □

Proof of Theorem 2. The proof proceeds by extending the idea of variance decomposition in the finite case to the limit as T tends to infinity. For every fixed T, as before, one defines

$$U(T) := F_1(X_1)^{Z_1} F_2(X_2)^{Z_2} \dots F_T(X_T)^{Z_T} \tag{17}$$

where $X = (X_1, X_2, \dots, X_T)$ is a random vector such that marginally $X_1 \sim G_1, \dots, X_T \sim G_T$ and is independent of $Z(T) = (Z_1, \dots, Z_T)$, which is a multinomial with T cells and equal probability of success in each cell.

Now, the difference here with the finite case is that $U(T)$ is no longer uniform. However, it follows from the condition (1) that $U(T)$ converges in distribution (or, weak convergence, see Billingsley [1]), as T tends to infinity to U, which is a Uni$(0, 1)$ random variable. However, one can still proceed with a variance decomposition as in the last proof. Notice that equality (10) continues to hold for each finite T when we define H in (7) by replacing Z by $Z(T)$ and U by $U(T)$.

The equality in (14) also holds, but Y_i does not follow F_i since $U(T)$ is not uniform. However, from (13), one can still write the following.

$$Var(H(Z(T), U(T))) \leq \frac{1}{T}\sum_{i=1}^{T} E\left(F_i^{-1}(U(T)\right)^2.$$

Combining the observations above, we get that for every fixed T, one has

$$E\left[\frac{1}{T}\sum_{i=1}^{T}\left(F_i^{-1}(U(T)\right)^2\right] \geq \frac{1}{T}\sum_{i=1}^{T} Var(G_i). \tag{18}$$

Now $U(T)$ converges to U in distribution. Define the sequence of functions

$$\theta_T(u) = \frac{1}{T}\sum_{i=1}^{T}\left(F_i^{-1}(u)\right)^2, \quad u \in (0,1).$$

Then, each θ_T is continuous and bounded by our assumptions. Also, as T tends to infinity, θ_T converges to θ uniformly. Thus, from the standard theory of weak convergence, it follows that $\lim_{T\to\infty} E\theta_T(U(T))$ exists and is given by $E\theta(U)$ where U is Uni$(0,1)$. In particular, from (18), it follows that

$$E\theta(U) \geq \limsup_{T\to\infty} \frac{1}{T}\sum_{i=1}^{T} Var(G_i). \tag{19}$$

Assume, for now, that each F_i has the same zero mean. Since θ_T converges to θ pointwise and remain uniformly bounded, it follows from the Dominated Convergence Theorem that

$$E\theta(U) = \lim_{T\to\infty} E\theta_T(U) = \frac{1}{T}\sum_{i=1}^{T} Var(F_i), \tag{20}$$

since $F_i^{-1}(U)$ has the distribution F_i and we have assumed its mean to be zero. Combining (19) and (20) we have proved the stated proposition.

Finally, the mean zero assumption on F_i's can be easily removed by suitable translations as was shown in the proof of the last proposition. □

Acknowledgment

Many thanks to Tilmann Gneiting for introducing me to the problem and for the subsequent comments, discussion, and encouragement.

References

[1] Billingsley, P. (1999) *Convergence of Probability Measures*. John Wiley & Sons, New York, NY.

[2] Bröcker, J. (2008) Decompositions of proper scores. Technical report. Available at `http://arxiv.org/abs/0806.0813`.

[3] Duffie, D. and Pan, J. (1997) An overview of value at risk. *J. Deriv.* **4**, pp. 7–49.

[4] Dawid, A. P. (1984) Statistical theory: the prequential approach (with discussion). *J. R. Statist. Soc. A* **147**, pp. 278–292.

[5] Dawid, A. P. and Vovk, V. G. (1999) Prequential probability: principles and properties. *Bernoulli* **5**, pp. 125–162.

[6] DeGroot, M. H. and Fienberg, S. E. (1983) Assessing probability assessors: calibration and refinement. In *Statistical Decision Theory and Related Topics III* vol. **1** (eds. S. S. Gupta and J. O. Berger), pp. 291–314. New York: Academic Press.

[7] Diebold, F. X., Gunther, T. A., and Tay, A. S. (1998) Evaluating density forecasts with applications to financial risk management. *Int. Econ. Rev.* **39**, pp. 863–883.

[8] Gneiting, T. (2008) Editorial: Probabilistic forecasting. *J. R. Statist. Soc. A* **171**, pp. 319–321.

[9] Gneiting, T., Balabdaoui, F., and Raftery, A. (2007) Probabilistic forecasts, calibration and sharpness. *J. R. Statist. Soc. B* **69**, part 2, pp. 243 – 268.

[10] Gneiting, T., and Raftery, A. (2005) Weather forecasting with ensemble methods. *Science* **310**, pp. 248 – 249.

[11] Gneiting, T., Stanberry, L. I., Grimit, E. P., Held, L., and Johnson, N. A. (2008) Assessing probabilistic forecasts of multivariate quantities, with an application to ensemble predictions of surface winds. *TEST* **17** (2), pp. 211–235.

[12] Granger, C. W. J. (2006) Some thoughts on the future of forecasting. *Oxf. Bull. Econ. Statist.* **67S**, pp. 707–711.
[13] Hamill, T. M. (2001) Interpretation of rank histograms for verifying ensemble forecasts. *Mnthly. Weath. Rev.* **129**, 550–560.
[14] Shafer, G. and Vovk, V. (2001) *Probability and Finance: It's only a Game!* New York: Wiley.
[15] Vovk, V. and Shafer, G. (2005) Good randomized sequential probability forecasting is always possible. *J. R. Statist. Soc. B* **67**, 747–763.

Department of Mathematics, University of Washington, Seattle, WA 98195
E-mail address: soumik@u.washington.edu